Asset Pricing and Earnings Fluctuations in a Dynamic Corporate Economy

William Gordon Ritter

Harvard University Department of Physics

17 Oxford St., Cambridge, MA 02138

Email: ritter@fas.harvard.edu

Phone: 857-919-2572, Fax: 617-495-0416

Abstract

We give a new predictive mathematical model for macroeconomics, which deals specifically with asset prices and earnings fluctuations, in the presence of a dynamic economy involving mergers, acquisitions, and hostile takeovers. Consider a model economy with a large number of corporations $C_1, C_2, \ldots, C_n$ of different sizes. We ascribe a degree of randomness to the event that any particular pair of corporations C_i, C_j might undergo a merger, with probability matrix p_{ij}. Previous random-graph models set p_{ij} equal to a constant, while in a real-world economy, p_{ij} is a complicated function of a large number of variables. We combine techniques of artificial intelligence and statistical physics to define a general class of mathematical models which, after being trained with past market data, give numerical predictions for certain quantities of interest including asset prices, earnings fluctuations, and merger/acquisition likelihood. These new models might reasonably be called "cluster-size models." They partially capture the complicated dependence of p_{ij} on economic factors, and generate usable predictions.

Key words: Cont-Bouchaud model, asset pricing, macroeconomic dynamics,

econophysics

PACS: 89.65.Gh

1 Introduction

A number of recent papers in theoretical macroeconomics have explored the utility of mathematical models coming from statistical physics. These theories typically model *agents* as vertices in a graph. Links between agents typically represent some form of coordinated behavior, in which clusters of agents have similar action on the market. The simplest and most successful model along these lines is the percolation model of Cont and Bouchaud [1], which entails a random communication structure between agents. Given two agents i and j, denote by p_{ij} the probability that they are linked (*linked* means these two agents have the same belief of the future movements of the asset price, and are similarly bullish or bearish). In order to obtain a mathematically tractable model, Cont and Bouchaud assume that $p_{ij} = p$ is a constant, independent of i and j. Thus the communication structure is a random graph, with connected components assumed to have the same action on the market. Such clusters are known to form in real financial markets, and they contribute to *herd behavior*, a phenomenon which is well-documented but not completely understood by theorists. However in real financial markets, clusters certainly do not form randomly; we therefore define in the present paper a model which attempts to capture this non-randomness.

In the Cont-Bouchaud model, at every time interval, each cluster randomly selects to buy (probability a), sell (probability a) or sleep (probability $1-2a$). The parameter a lies in the range $0 \leq a \leq 0.5$ and is called the *activity* of the

agents. The asset pricing structure of the Cont-Bouchaud model is very simple: aggregate excess demand, i.e. the sum of all orders, is the sole driving force of the asset price. Excess buying drives the price up, while excess selling drives it down. The price adjustment is therefore given by a price impact function

$$\ln P(t+1) - \ln P(t) = \frac{1}{b} \sum_{\text{agents } i} D_i(t), \tag{1}$$

where $P(t)$ denotes the price at time t, $D_i(t)$ is the demand of agent i in period t for the asset, and $b > 0$ is a liquidity parameter describing how much excess demand is necessary to move the asset price by one unit. The sign of $D_i(t)$ indicates buying or selling. In the Cont-Bouchaud model, $D_i(t)$ does not depend upon i. For each discrete time t, and for all i,

$$D_i(t) = \begin{cases} +1 & \text{with probability } a, \\ -1 & \text{with probability } a, \text{ and} \\ 0 & \text{otherwise} \end{cases}$$

This structure has been shown to mimic actual asset price dynamics quite closely [2].

Since a link is allowed to form between any pair of investors, the CB model is equivalently described as mean field theory of infinite-range bond percolation [1]. There is a critical value of p, denoted p_c and known as the *percolation threshold*, at which an infinite-range cluster can be formed. The change in log-price is proportional to $[\sum_{s\geq 1} s(n_s^+ - n_s^-)]$ where n_s^+ is the number of buying clusters of size s, and n_s^- is defined analogously for selling clusters. As discussed on p. 157 of [3], this number should be related to $\partial/\partial Q|_{Q=1}$ of the one-point function of the Q-state Potts model, although this connection has

[1] See Section 8.4 of [3] for background.

not been explored in the literature.

In real financial markets, p_{ij} is very far from being constant; in fact p_{ij} is a complicated function of a large number of variables[2]. The principal goal of the present paper is to give a new, predictive model for macroeconomic dynamics involving agents for which p_{ij} is not constant. A cornerstone of this new model, discussed in Sec. 3.3, is a procedure which effectively determines p_{ij} from past market data.

Although there does exist an analog of p_{ij} in the new model, and the new model does apply to asset pricing, it is not a direct generalization of the Cont-Bouchaud model, and is substantially more detailed. In our model, an individual site now represents a single unit of wealth. Agents in our model are firms (such as corporations or mutual funds) represented by clusters of sites. The size of the cluster is a crucially important variable representing the future earnings of the firm. The expectation of future earnings (EFE) may of course fluctuate as a function of time, and the amount of these fluctuations and which firms experience them are, like p_{ij}, complicated functions of a large number of variables.

In a real-world economy, fluctuations in asset prices are partially the result of rational investors following sound economic principles and seeking to optimize monetary gain, and partially the result of other factors, psychological or environmental, which by assumption do not lend themselves to scientific analysis and, for modeling purposes, may be viewed as random noise. The beauty of the new model introduced here lies in the fact that it works the same way. Fluctuations in the EFE of a firm, as well as the probabilities of mergers and

[2] We defer a discussion of which variables might be involved to Section 3.3.

acquisitions, are determined by a process which simultaneously incorporates both predictable economic trends, and the effect of randomness.

Monte Carlo time is a measure of efficiency of the simulation of a complex system; however in this paper, we take the approach that a Monte Carlo simulation of a lattice statistical system is a particular type of cellular automaton, and in simulations of dynamical systems using cellular automata, the "time" parameterizing the sequence of steps which update the cellular automaton is often identified with physical time. This identification of Monte Carlo time and physical time is key in what follows.

By way of introduction, we illustrate the simplest possible example of an interaction between two clusters of the same size. In later sections, we of course wish to generalize, allowing a large number of clusters of arbitrary sizes, and a much more detailed analysis of their interactions.

Imagine a hypothetical fledgling economy, in which there are a large number of consumers who demand a particular asset, but no producers of this asset. This is represented by a large two-dimensional lattice, with a minus sign on each lattice site to denote an absence of producers [3]. It will turn out that the dimension of the lattice must always be at least two in order to be useful to us; we defer a discussion of this point.

Since there is a demand without a corresponding supply, the system is not in equilibrium, and so two firms are created with the intent of producing this asset. Initially, the firms are comparable in size and difficult to distinguish in other respects, so analysts estimate the future earnings of each firm in the

[3] In physics, this is called the ground state of the Ising model

next quarter to be 4, in some appropriate units.

One lattice site represents one unit of expected future earnings, so each of these firms is denoted by the presence of a 2×2 matrix of plus signs occupying four lattice sites.

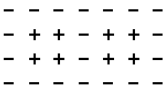

In our model, distance between connected regions in the lattice plays the role of likelihood for a merger. If a horizontal merger between the two firms in this hypothetical example is considered likely, we place them at a distance of only one lattice site from each other as in the figure. These simple binary variables which take their values in the set $\{+, -\}$ are called *spins* due to an analogy with physics.

We define a *link* to be an unordered pair of adjacent lattice sites. For example, if the lattice is the 2d integral lattice $\mathbb{Z} \times \mathbb{Z}$, then $b = \{(i, j), (i+1, j)\}$ is a link for all $i, j \in \mathbb{Z}$. We define the sign of the link to be the product of the signs at its two endpoints; so $(i, j) = +1, (i+1, j) = -1$ represents a negatively signed link.

We define the "energy" of the configuration to be

$$E = N_{+-} \tag{2}$$

where N_{+-} denotes the total number of negatively signed links. It follows that the system with two firms depicted above has energy $E = 16$.

Now, suppose that the firms decide for various reasons to undertake a horizontal merger. In statistical mechanics, a system will tend to relax to a state of lower energy. There are at least two lower energy configurations which achieve

the merger,

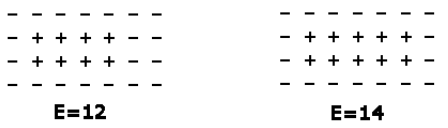

one of which involves region A moving to become adjacent to region B ($E = 12$), and the other involving the creation of two additional $+1$ lattice sites between regions A and B, in order to join the regions ($E = 14$). Since after the merger, the conglomerate will have access to all of the original clients of *both* firms, the expected future earnings of this new firm will be 8 or 10 units, depending on the details of the merger as in the previous figure.

A process is called *energetically favorable* if it leads to a state of lower energy. It is now clear that any positive multiple of (2) defines an energy operator in such a way that the merger of two nearby firms is an energetically favorable process. The toy model involving two firms is too simple to describe an actual economy, but the lattice described here forms the basis for a classic model of statistical physics (the Ising model), and the energy (2) is one term in the full energy operator of that theory; see equation (4). In subsequent sections, we will develop the full model which entails this as a special case. We will construct a cellular automaton which converges to equilibrium, and describe one method of making this into a predictive simulation of a large-scale economy.

A key point is the following. In the $E = 14$ example above, two 'spins' are flipped from $-$ to $+$ in the region between the two firms, as was necessary for the regions to merge. But a predictive model must tell us where to flip the next spin. In a large lattice with many firms, how would the site of the next flip be chosen? In Sec. 3.3, we outline a procedure for making this choice, involving a Markov model which learns from past behavior of the economy.

We give a simple example to illustrate the behavior of this Markov model. Suppose now that there are three firms A, B, and C, and as in the previous example, we give them all EFE $= 4$. However, these firms differ in other characteristics, and in particular, the book values are given by $v_A = 1, v_B = 2$, and $v_C = 10$. Also suppose that this hypothetical economy has a long history, in which economists observe that, in each past case involving one very wealthy firm (high book value) and a number of other firms with relatively small book value, the wealthy firm always acquired the least expensive of its competitors. It is then reasonable to expect, given this trend, that C will acquire A. Of course, this trend behavior is not completely certain (trends are sometimes broken), so the model should be probabilistic, while still retaining sensitivity to past economic trends.

All of the above goals are simultaneously accomplished by the following strategy: choose the location of the next spin flip to be between A and C (because $A + C$ is the most likely merger based on past trends), but flip the spin conditionally with a probability distribution that depends on the energy.

2 The Ising Model and Monte Carlo Iteration

In this section we introduce important terminology and define the basic model with which the rest of the paper is concerned. The model will be generalized and applied to economics in Section 3.

2.1 The Ising Model

The Ising model was originally introduced as a model for ferromagnetism, and has since been generalized and applied to many different fields, ranging from finance to biology. See [3] for a review. The system is based on a lattice with $\mathcal{N}$ sites in d dimensions, with a binary "spin" variable $\sigma_i \in \{\pm 1\}$ at each site i. While our examples are given in $d = 2$ to simplify the diagrams, for a realistic simulation it is likely that $d > 2$ is necessary. This technical issue is discussed in an appendix.

The model is defined by its energy function, which in the absence of an external magnetic field, is taken to be

$$E = -J \sum_{\langle ij \rangle} \sigma_i \sigma_j \tag{3}$$

where $J > 0$ is a coupling that determines how strongly nearby spins tend to align.

As in the physics literature, $\sum_{\langle ij \rangle}$ denotes a sum over all pairs (i, j) of sites which are nearest-neighbors in the lattice. We let N_{+-} denote the number of unlike nearest-neighbor pairs in the lattice, and let N_{++} or N_{--} denote the number of like nearest-neighbor pairs which are $+$ or $-$, respectively. It is not hard to see that (3) can be rewritten as

$$E = -J(N_{++} + N_{--} - N_{+-}) = -J(\mathcal{N} - 2N_{+-}) \tag{4}$$

Therefore, the energy consists of a ground state contribution $-J\mathcal{N}$ plus the positive constant $2J$ times (2).

A lattice has the defining property that each site has the same number of

nearest-neighbors. This number is denoted by q and is called the *coordination number* of the lattice. If we wish to embed the lattice in a plane using regular polygons, there are exactly three such tessellations of the plane:

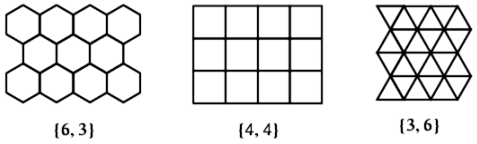

These generate regular lattices with $q = 3, 4$, and 6 respectively.

The thermodynamic behavior of a canonical ensemble is governed by the Boltzmann weight,

$$p(\sigma) \propto e^{-E(\sigma)/kT}$$

where k is Boltzmann's constant, and T is the temperature. The combination $1/kT$ occurs frequently, and so it is given the name β. This β has no relation to the coefficient in the linear Capital Asset Pricing Model, which measures the volatility of a security relative to its asset class. We will work in units for which Boltzmann's constant $k = 1$, in which case β is inverse temperature.

The ensemble average $\langle M \rangle$ of an observable M at temperature T is given by

$$\langle M \rangle = Z^{-1} \sum_{\sigma} M(\sigma) e^{-\beta E(\sigma)}$$

where $Z = \sum_{\sigma} e^{-\beta E(\sigma)}$ is the *partition function*. The notation $\sum_{\sigma}$ denotes a sum over the $2^{\mathcal{N}}$ possible states. It is not possible to calculate such an average by enumeration of all possible states, since an $L \times L$ square lattice entails 2^{L^2} states. However, ensemble averages can be accurately estimated with Monte Carlo methods.

A naive Monte Carlo estimate would draw configurations $\sigma^{(1)}, \ldots, \sigma^{(n)}$ from

the uniform distribution and calculate

$$\langle M \rangle \approx \left(\frac{1}{n} \sum_{j=1}^{n} M(\sigma^{(j)}) e^{-\beta E(\sigma^{(j)})} \right) \times \left(\frac{1}{n} \sum_{j=1}^{n} e^{-\beta E(\sigma^{(j)})} \right)^{-1}$$

The energy E is an extensive quantity, which means there will be huge fluctuations in $e^{-\beta E}$, and naive Monte Carlo is not a useful method for this model.

However, if $\sigma^{(1)}, \ldots, \sigma^{(n)}$ are drawn from the Boltzmann distribution rather than the uniform distribution, the naive Monte Carlo estimate is replaced with

$$\langle M \rangle \approx \frac{1}{n} \sum_{j=1}^{n} M(\sigma^{(j)}) \,. \tag{5}$$

Assuming that we have an efficient way to generate a sequence of Monte Carlo states that satisfy the Boltzmann distribution, then (5) is a remarkable improvement. We will now describe how to generate a sequence of Monte Carlo states according to the Boltzmann distribution; this is known as the *Metropolis algorithm*. It was first published in [4], and subsequently generalized and applied to many fields. For example, a similar algorithm [5] gives a Monte Carlo calculation of the ground state wave function in quantum mechanics. See also [6] and references therein.

2.2 The Metropolis Algorithm

The Metropolis algorithm was proposed as an algorithm to simulate the evolution of a system in a heat bath towards thermal equilibrium. From a given state i of energy E_i, generate a new state j of energy E_j by a small perturbation. If the proposed new state j has smaller energy than the initial state, then make j the new current state; otherwise accept state j conditionally, with

probability

$$A_{ij} = \exp\Big(-\beta(E_j - E_i)\Big)$$

After a number of iterations, we would expect to visit states of different energies according to the canonical distribution. This can be rigorously justified by showing that the sequence of states visited by the algorithm forms an ergodic Markov chain with the canonical distribution as a stationary distribution.

In more detail, the Metropolis algorithm is a stationary Markov chain on the state space, which generates a new state x' from the current state x as follows:

(1) Select a candidate state x^*, in which all components other than the kth are the same as in x, while x_k^* is chosen at random from a proposal distribution $S_k(x, x_k^*)$.
(2) Accept the candidate state with probability $A(x, x^*)$; otherwise, reject it and retain the current state.

We will always assume the proposal distribution is *symmetric*, in the sense that $S_k(x, x_k^*) = S_k(x^*, x_k)$ whenever $x_i^* = x_i$ for all $i \neq k$. Clearly, a sufficient condition for the Markov chain to be ergodic is that $S_k(x, x_k')$ is nonzero for all $x_k' \neq x$ and $P(x)$ is nonzero for all x. *Detailed balance* is a sufficient (but not necessary) condition for existence of equilibria, which says

$$A(x, x')P(x) = A(x', x)P(x')\,.$$

Whether or not the limit of the Metropolis algorithm satisfies detailed balance depends on the specific acceptance function.

For the pure Ising model, $S_i(x_i, x_i')$ is given by the delta function $\delta(x_i' - x_i)$, which means the candidate state is always obtained by flipping spin i to its opposite value. The acceptance function originally chosen by Metropolis *et al.*

is

$$A(x, x') = \min(1, P(x')/P(x)) \tag{6}$$

where $P(x)$ is the canonical Boltzmann distribution with respect to some energy function. For this choice of $P(x)$,

$$A(x, x') = \min(1, e^{-\beta \Delta E}), \quad \Delta E = E' - E \tag{7}$$

Suppose that state x' is obtained from state x by flipping the spin at location i. Then

$$\frac{1}{J}E = -\sum_{j \in nn(i)} \sigma_i \sigma_j \; - \sum_{\substack{\text{links } (k,l) \text{ not} \\ \text{involving i}}} \sigma_k \sigma_l$$

$$\frac{1}{J}E' = -\sum_{j \in nn(i)} (-\sigma_i)\sigma_j \; - \sum_{\substack{\text{links } (k,l) \text{ not} \\ \text{involving i}}} \sigma_k \sigma_l$$

where $nn(i)$ denotes the set of nearest-neighbors of site i. Therefore,

$$\Delta E = E' - E = 2J \sum_{j \in nn(i)} \sigma_i \sigma_j \tag{8}$$

In particular, evaluating the acceptance probability for any one spin flip only requires calculation of q links.

The probability of accepting state x' is therefore found to be

$$A(x, x') = \min\Big(1, \exp\Big(-2K\sigma_i \sum_{j \in nn(i)} \sigma_j\Big)\Big) \tag{9}$$

with $K = \beta J$. This is perhaps the simplest acceptance function for which detailed balance holds. However, a large class of acceptance functions are known for which the Metropolis algorithm satisfies detailed balance. Aside from (9), the most commonly used one is the Boltzmann acceptance function

$$A(x, x') = P(x')/(P(x) + P(x'))\,, \tag{10}$$

which, for distributions defined by an energy function, yields

$$A(x, x') = \frac{e^{-\beta E'}}{e^{-\beta E} + e^{-\beta E'}} = \frac{1}{1 + e^{\beta \Delta E}} \,.$$

where ΔE was calculated in (8). It is easy to show that both (6) and (10) lead to systems satisfying detailed balance.

2.3 Cellular Automata

Generally speaking, a *cellular automaton (CA)* consists of a d-dimensional lattice of "cells," each in one of a finite number of states. The grid can be in any finite number of dimensions. Time is also discrete, and the state of a cell at time t is a function of the state of a finite number of cells called the "neighborhood" at time $t-1$. In the traditional interpretation, this function is not considered to contain any random variables, and thus the time evolution of the system is deterministic. However, an important generalization (and the one which we will use) is the case in which the state-update function is stochastic; such a system is termed a *probabilistic cellular automaton (PCA)*.

Every cell is assumed to have the same rule for updating, based on the values of neighboring cells. Each successive application of the rules to the whole grid is called a "generation." The number of updates in a Monte Carlo simulation is sometimes referred to as *Monte Carlo time*. However, units of Monte Carlo time simply label generations in the CA, and in general, cellular automata provide a natural language in which to describe simulations of lattice systems.

The Ising model with energy (3), evolved forward in time by the Metropolis algorithm (defined with the acceptance function (9)), clearly forms a probabilistic cellular automaton. Henceforth, we refer to this PCA as the *Metropolis-*

Ising system.

3 Application to Economics

3.1 Assumptions

To apply the Metropolis-Ising cellular automaton described in the previous section to economics, we make the following assumptions.

(1) *(Economic interpretation)* Lattice sites represent units of wealth. A connected region of $\mathcal{E}$ spin-up (+) sites represents a company with expected future earnings equal to $\mathcal{E}$. Spin-down (−) sites represent vacancies.

(2) *(Initial state)* The lattice is initialized with a state in which each currently existing company in the economy under consideration[4] is denoted by a connected region of appropriate size, proportional to that firm's expected future earnings. Additionally, the distance $d(C_i, C_j)$ between companies C_i and C_j is proportional to a current estimate of merger likelihood.

(3) *(Time evolution)* The predicted time evolution after the initial state is given in discrete steps, with state x_{i+1} determined from state x_i by conditionally flipping the spin at lattice site τ_i, with probability determined

[4] The global economy is certainly too large to effectively apply this model. By the "economy under consideration" we mean a particular sector of the economy which, from the point of view of mergers, and acquisitions, may be considered as approximately independent of other sectors. This assumption does *not* entail zero trade with the other sectors.

by the Metropolis function

$$\min(1, e^{-\beta\Delta E}),$$

or a different acceptance function leading to detailed balance, such as

$$(1 + e^{\beta\Delta E})^{-1}.$$

Here, β denotes the inverse temperature; an appropriate value of the temperature for use in economics is determined in Section 3.2. The location of the site τ_i is also crucial, and is determined by an artificial intelligence algorithm (Section 3.3) which learns from past market behavior.

(4) *(Nearest-neighbor interaction)* The energy function E includes a nearest-neighbor interaction $-J \sum_{\langle ij \rangle} \sigma_i \sigma_j$.

Axiom 1 as it stands is appropriate for the intended application to earnings fluctuations. It will be modified in Section 3.4 when we consider the application of our methods to asset pricing.

Axiom 4, the nearest-neighbor interaction, needs some explanation. Based on Axiom 1, connected regions of spin-up (+) sites represent firms, which are the agents in this model; it must be statistically possible for mergers to occur and for small clusters to be randomly created, and this is true of the energy function defined in Axiom 4. A more sophisticated justification for Axiom 4 is that the Ising Hamiltonian is the simplest of a class of lattice interactions which have a phase transition and associated critical point. This is important because it means that there is an adjustable parameter ($K = \beta J$) in the theory in terms of which the correlation length is an unbounded function [5]. Since the

[5] More precisely, it is unbounded for an infinite lattice. For a finite lattice, it can grow to at most the size of the lattice. See Ref [3] for a further discussion.

correlation length is the typical size of clusters, it follows that without an Ising-like phase transition, there might be no temperature for which we have the typical cluster size observed in the economy. That is to say, equation (11) might have no solution. The energy function $-J\sum_{\langle ij\rangle}\sigma_i\sigma_j$ is merely the simplest in a large class of known energy functions which guarantee existence of a phase transition, and hence existence of solutions to (11).

In real-world financial markets, the expected future earnings referred to in Axiom 1 tend to fluctuate partially according to predictable economic trends, and partially due to unpredictable events, which can be viewed as randomly occurring. The model proposed in this paper is appropriate in this sense, as it is a stochastic model in which the sizes of connected regions can experience random fluctuations. However, like the merger probabilities, these fluctuations in future earnings are also partially determined by intelligently modeling past economic trends; see Section 3.3 for further discussion of this point.

Just as the original Ising model is only an approximate description of a ferromagnet, there are a number of reasons why the Metropolis-Ising model described here can only be a rough approximation to a real-world economy. When two connected regions in the lattice merge, the size of the new region formed is (at least for a few units of Monte Carlo time) roughly the algebraic sum of the sizes of the two constituent regions. This corresponds in microeconomic theory to the assumption that under a merger or acquisition, the expected future earnings of the conglomerate is given approximately by the algebraic sum of the EFE's of the two merging firms, which may or may not be the case depending on the details of the merger.

However, in spite of these approximations, the model, with its given macroeconomic interpretation, has a number of advantages. One advantage is that the Metropolis algorithm, which defines the Ising cellular automaton, was originally invented in order to facilitate numerical simulation of the Ising model in various dimensions. In our interpretation, Monte Carlo time corresponds to real time in the economy, but the natural amenability to computer simulation remains. [6]

3.2 The Critical Temperature

We now discuss the notion of *temperature*, one of the free parameters of the Ising model, and determine the correct "temperature range" for the current economy.

At temperatures well above the critical temperatures, the spin arrangement converges in Monte Carlo time to a nearly random arrangement, independent of the starting state, and fluctuates quickly. In physics terminology, above the

[6] There is an extensive literature on the "critical slowing-down" effect for Metropolis-type simulations. Stated precisely, the problem is that the autocorrelation time diverges as a power-law with increasing system size, near the critical point. In other words, the computing time required to generate an independent configuration increases superlinearly with the system volume. Fortunately, in the economic application described here, we are not concerned with generating independent configurations or approaching equilibrium; our sole concern is that the system continue to correctly model the economy, and this is guaranteed (within the approximations we make) by the intelligent choice of spin update sites through the use of our Markov model. Thus critical slowing down is <u>not</u> a problem!

critical temperature there is a single thermodynamic state with zero magnetization. This corresponds to a nightmarish economy in which firms of any size form, merge, split, and terminate frequently and randomly; the number of companies in existence tomorrow has no correlation with the number today. These values of the temperature do not describe any real-world economy.

Below the critical temperature, there are two thermodynamic states (the "up spin" state with positive magnetization and the negative magnetization "down spin" state) and the system stays in one or the other depending on how the spins are initialized. This corresponds to a completely socialist economy in which either there are no firms, with all industry controlled by the government, or all corporations have merged into a single firm.

The actual economies of most countries fall somewhere between these two limiting cases. Near the critical temperature in the Metropolis-Ising model we expect large clusters of spins with the same orientation, which fluctuate, but very slowly. Near-critical temperature corresponds, under our interpretation of the clusters as companies, to the behavior of a large free-market capitalist system such as the U.S. economy or global economy. [7]

The typical size of clusters is called the *correlation length* ξ, which is maximal at the critical temperature for a finite system, and which diverges at T_c for an infinite system. Section 3.1 detailed a procedure for predicting the time evolution of an economy, within the bounds of a simple lattice model. This predictive procedure is defined in terms of a parameter β = inverse tempera-

[7] This shows in particular that the one-dimensional Ising model is useless for the particular macroeconomic interpretation advocated here, as it has no phase transition.

ture, so to use the model, we need a numerical value for the temperature. In the predictive time evolution, we know the state of the economy now, which (when drawn on a lattice) contains a number of clusters. Assume that we have calculated the average cluster size of this initial state, and denoted this number by ξ_0. This is also called the correlation length of that state. Near criticality in the $d = 2$ Ising model, the correlation length displays the power-law behavior $\xi \approx (4|K_c - K|)^{-1}$, where $K = \beta J = J/T$ (see [8], p.387).

These observations indicate that a good choice for the temperature is $T_0 = J/K_0$, with K_0 defined as either of the two solutions to the equation

$$\frac{1}{4|K_0 - K_c|} \;=\; \xi_0 \,. \tag{11}$$

In words, eqn. (11) instructs us to determine the temperature needed for the Metropolis update procedure based on the apparent temperature of the initial state. The latter is calculated from (an approximation to) its correlation length.

The value of K at criticality is determined by a simple self-duality argument ([8], p. 383). Its numerical value is

$$K_c = \frac{1}{2}\sinh^{-1} 1 = \frac{1}{2}\ln(\sqrt{2}+1) \simeq 0.4407$$

If we work at subcritical temperature, $T_0 < T_c$, we then have explicitly

$$K_0 = \frac{1}{4\xi_0} + 0.4407$$

Both the Metropolis acceptance function (9) and the Boltzmann acceptance function (10) depend on β and J only through the combination $K = \beta J$.

3.3 The Markov Model

This is the most important section in the paper, as here we discuss how to use the Metropolis-Ising model to obtain predictions for real-world financial markets. As we shall see, a Markov-model determination of the update location, τ_i, is the key.

The Metropolis-Ising model, as presented above, can certainly be considered as a qualitative model for macroeconomics. It allows for an initial configuration to evolve forward in time, while tending to a state of thermodynamic equilibrium. The evolution is governed by an energy function E, together with one adjustable parameter $K = \beta J$ which depends on the coupling and the temperature, and this energy function has the property that mergers of economically aligned firms tend to give a state of lower energy, and for that reason are statistically likely. Also likely (both in the economy and in our model) is the random formation of new companies with small market share. The model does not support spontaneous creation of very large firms, and neither does the economy. In summary, the properties of time evolution in this model all agree qualitatively with the properties of an arbitrary, generic large-scale corporate economy, but the model as it stands is too general to give specific properties of one particular economy.

The problem of over-generality of a model is common in physics. Einstein's equations for gravity admit thousands of known solutions, and only a fraction of these correspond to physically possible configurations.

For similar reasons, the Metropolis-Ising model will be of very limited use in modeling real-world economies unless it can be trained to take advantage of

the fact that history repeats itself. The world's financial markets are known to exhibit various cycles, or patterns (such as bubble economies) which are self-similar, in the sense that the $(n+1)$st bubble can be compared with the first n bubbles and certain similarities emerge.

We will show in this section that it is possible to program the Metropolis-Ising model with historical data in such a way that a particular pattern of corporate mergers or splitting which occurred in the past is more likely to occur in a similar pattern again. The technique is borrowed from artificial intelligence.

A *Markov model of order* n is a semi-random walk through a state space Σ. Given $i-1$ states, the i^{th} state is chosen based on a probability distribution over Σ which may depend on the previous n states, but not on states further back. A common choice for the probability distribution is described as follows. For simplicity, consider $n = 2$, so a second-order Markov model. Suppose that we have a sequence (x_i) of length $k \gg 2$ which represents behavior we wish to mimic. For each ordered pair $(a, b) \in \Sigma^2$, we define a histogram $p_{a,b}$ over Σ by $p_{a,b}(c) =$ conditional probability, given that a, b were the ith and $(i+1)$st elements in the sequence, that c is the $(i+2)$nd element. The process of determining $p_{a,b}$ for all $(a, b) \in \Sigma^2$ by iterating over the elements of (x_i) is called *training*. This clearly extends to general n. A typical application of this technique is to fill in a missing note in a Mozart piano concerto, by first training the Markov model on all of Mozart's other works.

More generally, let S, T, and U be any sets. One might have two sequences $(s_i) \subset S$ and $(t_i) \subset T$, and a third sequence f_i taking values in U, which depends on the first two in some complicated way that we wish to mimic. (Equivalently, there is a complicated function $f(i, s, t)$ on $\mathbb{N} \times S \times T$ which is

unknown; all that is known are its values $f(1, s_1, t_1), f(2, s_2, t_2), \ldots$) For each ordered pair $(a, b) \in S \times T$, define a histogram $p_{a,b}$ on U such that $p_{a,b}(u)$ equals the number of times $f_i = u$ given that $s_i = a$ and $t_i = b$. It is actually this more general kind of second-order Markov model which we will use. In our application, $S = T =$ the space of possible *key statistics* of a company (defined below), and $U =$ the space of lattice triangles. The histograms could each be normalized so that $\sum_{u \in U} p_{a,b}(u) = 1$, in which case the $p_{a,b}(u)$ may be viewed as probabilities.

The first step in applying this to the Metropolis-Ising model is that a computer must generate, from past financial data about the number of companies as a function of time (measured in discrete time-slices separated by δt) and their relative amounts of wealth, a sequence of Ising model configurations $x_1, x_2, \ldots, x_N$ in such a way that

(1) *(Single-site updates)* The state x_{i+1} differs from x_i by the flip of a single spin. Denote the location of this spin by τ_i.

(2) *(Historical Accuracy)* There exists a subsequence $x_{i_1}, x_{i_2}, \ldots$ such that x_{i_j} correctly represents a snapshot of the market at time $t_0 + j \cdot \delta t$.

(3) *(Present Accuracy)* The last configuration, x_N, represents the current state of the economy, so $x_{N+1}, x_{N+2}, \ldots$ are unknown, as are τ_N, τ_{N+1}, etc.

Condition 2 needs some explanation. We divide time into intervals of length δt; since earnings are often reported at the end of each quarter, it probably makes most sense for δt to correspond to one fiscal quarter. The amount of merger activity and earnings fluctuations for the various firms in an economy, per unit time, clearly changes over time. This activity is directly related to the

number of spin flips needed to propagate the lattice from one time interval to the next. From state x_{i_0}, corresponding to time t_0 (the earliest date for which data is available), suppose that n spin-flips are needed to reach the state of the economy at time $t_0+\delta t$. In this situation, our notation is such that $i_1 = i_0+n$.

It is now possible to train a Markov model or neural network using the sequence of histories $\{x_i\}$; we will describe the Markov model approach. What exactly should be the data sets (elements of Σ in our previous notation) on which to train the AI? The naive choice, to train on the Ising configurations themselves, will certainly be useless since the exact positions of the companies on the grid is arbitrary; the model is invariant under translations and rotations. A more sophisticated training approach is necessary.

To this end, we define the notion of *key statistics* for a company. By this we mean any collection of purely numerical data about the company which an analyst might consider when studying mergers involving that company, or when studying its earnings. These variables include, but are not limited to: stock valuation measures, profitability, return on equity, revenue, earnings growth, total cash, total debt, and many other indicators. Using more indicators might make the model's predictions more accurate, but too many variables will lead to an undesirable Markov model, having a large number of distinct histograms $p_{a,b}(u)$ which are all trivial. For a company C, we denote its key statistics by $\S(C)$; this is a vector in $\mathbb{R}^{\mathcal{K}}$ where $\mathcal{K}$ is the number of key variables.

If it occurs in a periodic or semi-periodic fashion that a company with key statistics $\S$ tends to merge with a second company (key statistics $\S'$) whenever all other market conditions are equal, then that is an economic trend. In what follows, we determine a Markov model which will certainly capture such a

trend if it exists. If there is no relation between key statistics and the prediction of a merger, then the problem is completely intractable and it would be impossible even in principle for a model to be developed; we do not take this view.

The Markov model we propose is an artificial intelligence technique which chooses the position of the next Metropolis update point intelligently, based on learning of past market behavior. By assumption, for all $i \geq 1$, configuration x_{i+1} differs from x_i by a single spin-flip; denote the site of this flip as before by τ_i. The fundamental problem is then to determine the locations of τ_N, τ_{N+1}, and so forth. Knowledge of the τ_j for $j \geq N$ tells us how to calculate the x_j as well, since x_j is determined from x_{j-1} by calculating the value of the acceptance function A and (probabilistically, according to the Boltzmann weight) flipping site τ_{j-1}.

To train the Markov model, for each i we identify the unique pair of companies (A_i, B_i) that minimizes the sum of the squares of the distances from τ_i, i.e. choose A_i, B_i to minimize

$$d(A_i, \tau_i)^2 + d(\tau_i, B_i)^2. \tag{12}$$

and order the pair (A_i, B_i) so that

$$d(A_i, \tau_i) < d(\tau_i, B_i) \tag{13}$$

We note that if τ_i is located at the boundary of a company C, this is directly related to fluctuations of the earnings potential of that company. In that situation, conditions (12) and (13) conspire to ensure that C will be the first element of the ordered pair that is chosen.

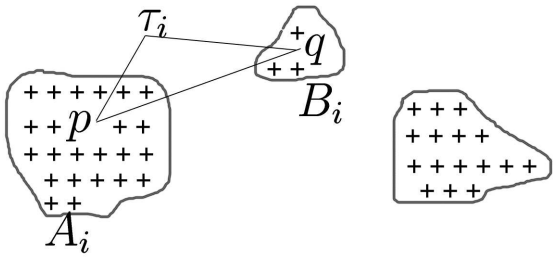

For each such pair of companies, we will train the Markov model on the ordered pair $(\S(A_i), \S(B_i))$. Since this is not as simple a Markov model as the Mozart missing-note example discussed above, it is nontrivial to define what we mean by *training*. Let p, q denote central points in regions A_i, B_i respectively. With τ_i, these form a triangle $\triangle pq\tau_i$. Since scaling the entire model is equivalent to a change of monetary units, we are only interested in the similarity class of this triangle, which may be encoded, for example, with two angles.

In the histogram notation from our generalized second-order Markov model, this entails setting

$$f(i, A_i, B_i) = \triangle p\, q\, \tau_i\,. \tag{14}$$

Two companies will never have exactly the same key statistics; when the Markov model is matching pairs $(\S(A_i), \S(B_i))$, it must use a slightly sophisticated matching function which compares the key statistics of two companies and returns "true" if they are similar, or lie in the same range [8].

Now suppose that we are in the situation outlined above: we have a sequence of Ising lattice configurations $x_1, x_2, \ldots, x_N$ which correctly represent the past history of the economy (within the accuracy of the lattice approximation).

[8] If certain key statistics are more indicative than others, the comparison function may also take the form of a linear combination of the various key statistics with some weights inserted.

Our task is to compute x_{N+1}, for which we need τ_N, the site of the next flip.

The new update site τ_N is chosen so that its position relative to the two closest companies mimics the relative positions of update sites in the past history of the economy. A simple geometric argument involving similar triangles supplies the correct notion of "relative position."

For each pair (A, B) of companies on the lattice, compute the key statistics $\S(A), \S(B)$, and look up the resulting ordered pair $(\S(A), \S(B))$ in the histogram f formed by the Markov model. Select the pair with the best match. By definition of the histogram (14), this returns a triangle which we will denote $\triangle_{AB}$. Denote by p' and q' the central points of A and B respectively. There is now a unique point τ_N such that $\triangle p'q'\tau_N$ is similar to $\triangle_{AB}$.

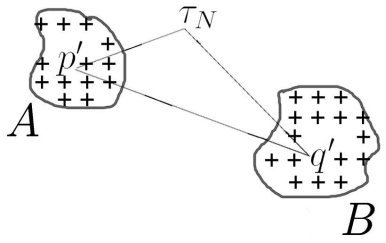

3.4 Asset Pricing

Our model is immediately applicable to asset pricing. In fact, the same asset price structure as in the Cont-Bouchaud model makes sense here as well, under the assumption that the amount of trading in this asset type by each cluster is proportional to the size of the cluster.

Previously, we wanted to estimate fluctuations in expected future earnings of a

firm [9], and so we worked under the approximation that, immediately following a merger, the expected future earnings of the conglomerate is the algebraic sum of the expected earnings of the two merging firms. It is always the case that when two clusters in a lattice cellular automaton merge, the size of the combined cluster is approximately the sum of the sizes of the constituents, for the first few generations after the clusters combine. Thus, the cluster-size models advocated in this paper can only be used to estimate *locally additive functions* [10]. The EFE was a locally additive function, and thus it made sense to identify its value with the number of sites forming the cluster, in some appropriate units of wealth. In asset pricing models (such as the Cont-Bouchaud model) which are driven by aggregate excess demand, the locally additive function we wish to estimate is now the demand of each agent for the asset, since in such models the price change is a linear function of the demand for the various agents, as in eqn. (15). It is worth noting that (11) gives a good estimate of the temperature to use, regardless of which locally additive function we are modeling.

Although not necessarily always true, it is a reasonable approximation that a conglomerate formed by a *recent* merger of firms A and B will carry out an amount of trade which is the algebraic sum of the two firms' trading before the merger. As before, we also assume that at every time interval, each cluster randomly selects to buy (probability a), sell (probability a) or sleep (probability $1 - 2a$), and aggregate excess demand is the driving force of the asset

[9] This is very important, since any reasonable stock valuation method would have to depend on EFE.

[10] "Locally" here means local in time, and refers to the fact that the functions fluctuate in time, but are additive immediately following the merger.

price.

The price adjustment is, as before, given by a price impact function

$$P(t+1) - P(t) = \frac{1}{b} \sum_{\text{clusters } i} D_i(t) \tag{15}$$

where $P(t)$ denotes the log price at time t, $D_i(t)$ is the demand of cluster i in period t, and $b > 0$ is a liquidity parameter describing how much excess demand is necessary to move the asset price by one unit. The demand $D_i(t)$ is defined as plus or minus the size of cluster i, with the sign determined as before: $+1$ with probability a, -1 with probability a, and zero otherwise. The departure from the Cont-Bouchaud model lies in the fact the structure of the clusters is no longer random. The number and sizes of the various clusters will fluctuate in an economically meaningful way, determined by artificial intelligence from past market behavior based on analysis of key statistics. These fluctuations generate important corresponding implications for the asset price.

Since $D_i(t)$ for an individual site can take values $-1, 0, +1$, this asset pricing system suggests that the appropriate statistical physics model for this application is, in fact, the *Blume-Capel model* [7,3]. The Blume-Capel model (sometimes called the *tricritical Ising model* due to the point in its phase diagram at which three lines of second-order phase transitions meet) is a generalization of the Ising model in which the spin variable on each lattice site can take three possible values $s = -1, 0, +1$. In physics, this is the classical version of a quantum spin-one magnet.

A typical configuration of the Blume-Capel model is a lattice with various connected regions of plus or minus sites, separated by sites r with $s(r) = 0$, which may be thought of as vacancies. The Blume-Capel model, and its

associated Metropolis cellular automaton (see Sec. 2.3), can be thought of as a model for mergers and acquisitions if we introduce assumptions similar to those of Sec. 3.1. It is now fruitful to let connected regions of plus or minus denote firms (or agents, abstractly), leaving the zeros as vacancies. The Blume-Capel model becomes an asset pricing model reminiscent of Cont and Bouchaud when we declare that, at each time interval, the plus sites buy the asset, the minus sites sell the asset, and the vacancies sleep. It becomes a *predictive* asset pricing model (which takes into account mergers/acquisitions and fluctuations in demand) when we train it with past market data using a Markov model of the type introduced in Sec. 3.3.

The Metropolis algorithm becomes slightly more complicated for the Blume-Capel model, because even after the update site is chosen, the spin can 'flip' to one of several possibilities. We assert that a reasonable Metropolis procedure, which satisfies detailed balance, is the following. Let $x_1^*, \ldots, x_k^*$ denote the different possible choices for the new state ($k = 2$ for the Blume-Capel model). If one of the x_i^* has lower energy than the current state, then switch to a new state which minimizes energy among the candidate states. Otherwise, choose the candidate state with the minimum $|\Delta E|$, and conditionally switch to that state with probability $\exp(-\beta \Delta E)$.

4 Conclusions

We have defined a class of new predictive mathematical models for *locally additive* economic functions (such as demand for an asset, or expected future earnings), which predicts the future behavior of that function based on past economic trends involving mergers, acquisitions, and fluctuations in the

function itself. This model builds upon previous models of statistical physics, Monte Carlo simulation, and artificial intelligence. Since demand for an asset is among the class of allowable functions, this immediately gives a new asset pricing model.

The Metropolis-Ising cellular automaton was seen to be a qualitative model for macroeconomic dynamics. It allows for an initial configuration to evolve in time, approaching a state of equilibrium, with evolution governed by an energy function E. The energy function we use is familiar from statistical physics, and has the property that mergers of economically aligned[11] firms lead to states of lower energy, and hence are statistically likely. Also likely (both in the economy and in our model) is the random formation of new companies with small market share and small short-term expected earnings. The model has one adjustable parameter $K = \beta J = J/kT$. The model does not support instantaneous creation of firms with very large market share, and neither does the economy. In summary, the properties of time evolution in this model agree qualitatively with the properties of an arbitrary, generic large-scale corporate economy.

A more important property of the model is the following. Two firms C_i and C_j are more likely to merge if the next Metropolis update point is chosen between the respective clusters. The position of this update point, which mostly determines the probability of a merger p_{ij} (see also Sec. 1), is chosen by a Markov model which effectively reproduces past economic trends. In real financial markets, and in this Markov model, p_{ij} is partially a complicated

[11] This means that the corresponding clusters are close together in units of lattice spacing.

function of a large number of variables which are called *key statistics*, and partially a function of purely random events. The same reasoning applies to predict the fluctuation in expected future earnings of a firm. The EFE is, like p_{ij}, a time-dependent pseudorandom variable which is partially determined by a complicated function of key statistics.

Our model is immediately applicable to asset pricing. In fact, the same asset price structure as in the Cont-Bouchaud model makes sense here as well, under the assumption that the amount of demand for this asset type by each cluster is proportional to the size of the cluster. In this application, it makes more sense to consider a slight generalization of the Ising lattice, which is known in physics as the Blume-Capel model [7]. The asset pricing model which this entails represents a dramatic generalization of the work of Cont and Bouchaud: the structure of the clusters is no longer random. The number and sizes of the various clusters in this new model will fluctuate in an economically meaningful way, with corresponding implications for the asset price.

The model is explicitly computable. At no point does it involve iteration over the $2^{\mathcal{N}}$ possible states of the Ising system; the acceptance function $A(x, x')$ is computed only from data involving the neighbors of a single lattice site.

Acknowledgements

The author is grateful to Daniel Jafferis for helpful discussions.

A Appendix: Geometric Limits

In our model, firms likely to merge must be in close proximity on the lattice. The following scenario may certainly arise: a certain sector of the economy is dominated by a single large firm, with a number N of smaller firms competing for the remaining market share.

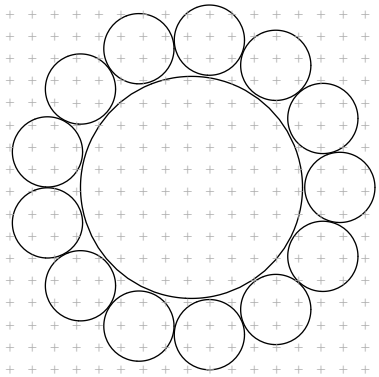

The large firm will acquire one or more of the small firms by the mechanisms described elsewhere in this paper. In what follows, we will assume that the regions are roughly spherical. This means that N small regions (radius r') must be in close proximity to one larger region (radius r) in the graph. In any dimension, there is a geometric limit on N, which is determined by a very simple sphere-packing bound, derived below. Of course, any such limit is an artificial restriction placed by the model, not a real restriction on the economy; fortunately, our estimates show that this problem is resolved by increasing the dimension d of the lattice.

For this estimate, we may approximate an integral lattice of rank d by the continuum limit $\mathbb{R}^d$. In $d = 2$, the situation depicted in the figure, the number of circles of radius r we can pack around a fixed circle of radius $R > r$ is

bounded as

$$N_2 \leq \pi \frac{R+r}{r},$$

since the sum of the diameters $2Nr$ should not exceed the circumference of a circle with radius $R+r$. A similar argument in $d=3$ yields

$$N_3 \leq 4\left(\frac{R+r}{r}\right)^2.$$

For concreteness, suppose the small companies are 1/100 the size of the large one. Then $(R+r)/r = 11$, and we have approximately $N_2 \leq 35,\ N_3 \leq 484$. Thus, in three dimensions, if a given firm is considering making an acquisition of a smaller firm with one-hundredth its wealth, there can be at most 484 such smaller firms. This may not be enough, but these bounds continue to grow in higher dimensions. We therefore assert that for sufficiently large dimension, the geometry contains sufficiently many degrees of freedom for this kind of economic modeling.

References

[1] R. Cont and J.P. Bouchaud, Macroeconomic Dynamics **4**, 170 (2000) `cond-mat/9712318`.

[2] Stauffer, D. "Percolation Models of Financial Market Dynamics," Adv. Complex Systems **4** no. 1, 19–27.

[3] Cardy, John L. *Scaling and renormalization in statistical physics*, Cambridge University Press, 1996.

[4] N. Metropolis *et al.*, J. Chem. Phys. **21** 1087 (1953)

[5] McMillan, W. L., Phys. Rev. A, **138**, 442 (1965)

[6] D. M. Ceperley, "Metropolis Methods for Quantum Monte Carlo Simulations," `physics/0306182`.

[7] M. Blume, Phys. Rev. **141**, 517 (1966).
H. W. Capel, Physica **32**, 966 (1966); **33**, 295 (1967); **37**, 423 (1967).

[8] Pathria, R. K. *Statistical Mechanics*, 2nd ed. Oxford; Boston: Butterworth-Heinemann, 1996.